

19990617 017

How Probability Theory Can Help Us Design Rule-Based Systems

Donald Bamber

Space & Naval Warfare Systems Center, San Diego
SPAWARSYSCEN D44215
53345 RYNE ROAD
SAN DIEGO CA 92152-7251

Abstract

This paper is concerned with finding improved methods of reasoning for use in rule-based systems that perform diagnosis in general and situation assessment in particular. It is argued that, because the rules used in rule-based systems typically have exceptions, the rules must be interpreted probabilistically. Thus, if a rule *If A then B* has exceptions, then what the rule really means is that the conditional probability of *B* given *A* is close to one. A rule-inference criterion that makes use of second-order probability concepts is advocated. Interestingly, this inference criterion is equivalent to some non-probabilistic inference criteria. The paper expounds a scheme for constructing situation-assessment systems that could be used as either decision aids or as reasoning components of computer generated forces.

1. Introduction

My goal in writing this paper is to give a non-technical summary of my work on the logic of "high conditional-probability assertions" and to discuss how that logic may be used in developing rule-based systems for situation assessment. In this paper, I have used terminology and given explanations that are designed to convey basic intuitions rather than to be technically precise. Those readers who would like to see a technically precise formulation of this work with full technical details may consult my papers [Bamber, 1996; Bamber, to appear].

As its title indicates, this paper discusses how probability theory may be used to help in the design of rule-based systems. The specific type of rule-based system that this paper is concerned with are those that perform diagnosis.

By *diagnosis* is meant the process of using known properties of an entity to deduce new (*i.e.*, previously unknown) properties of that entity. For example, in medical diagnosis, the entity to be diagnosed is a patient, the entity's known properties are symptoms, and the new properties to be deduced are medical conditions. As another example, in situation assessment, the entity to be diagnosed might be a particular ship, its known properties might be its speed and the frequency of its radar, and the new property to be deduced might be the ship's classification. Thus, diagnostic rule-based systems use known properties to deduce new properties. They are designed to ask questions of the type: *If an entity has properties A, B, C, etc., will it also have property D?*

(Diagnostic rule-based systems should be distinguished from *action evaluation/planning* rule-based systems that are designed to either evaluate the consequences/utility of actions or to plan actions. Such systems are beyond the scope of this paper.)

Because situation assessment is a form of diagnosis, rule-based systems for situation assessment are a potential application for the work presented in this paper. Such situation-assessment systems would be useful either (a) as a decision aid for leaders of real-life military forces

or (b) as a simulation of the situation assessment performed by computer generated forces.

2. Propositional Logic as a Method for Reasoning with Rules

In diagnostic rule-based systems, the rules take the form *If A then B*, where *A* and *B* each denote properties of the entity to be diagnosed. How ought we to reason with rules that have this form? The obvious answer would appear to be that we should use propositional logic (also known as *sentential logic*). After all propositional logic [Enderton, 1972, Chapter 1] is designed for reasoning with statements having various forms, among them being the form: *If A then B*. And it is widely regarded as being the correct logic to use for reasoning with such statements.

Terminology and notation. In the following, it will be convenient to say that the property *A* is true or holds true for an entity if the entity possesses property *A*. Given the properties *A* and *B*, define $\sim A$ to be the property that holds true if and only if *A* does not hold true. Define $A \& B$ to be the property that holds true if and only if *A* and *B* both hold true. Define $A \vee B$ to be the property that holds true if and only if either *A* or *B* holds true. The rule *If A then B* will be denoted $A \Rightarrow B$.

When propositional logic is used to reason with rules, the symbol " \Rightarrow " is interpreted as being the material-conditional connective. Recall that this connective is defined in such a way that $A \Rightarrow B$ is equivalent to $\sim A \vee B$. In other words, the rule *If A then B* is considered to be equivalent to *Either not A or B*.

(Later in this paper, an alternative interpretation of the rule *If A then B* will be considered. Because the rule will still be denoted $A \Rightarrow B$, it will be necessary to reinterpret the symbol " \Rightarrow ".)

2.1 A Problem Caused by Exceptions: Rare Errors

There is a problem with using propositional logic to reason with the rules that are found in rule-based systems. That problem is that propositional logic was designed to reason with statements that have no exceptions. The point of view taken in propositional logic is that, in order to justifiably assert a statement, that statement must always be true. One is not justified in asserting the statement *If A then B* if, when the property *A* holds, the property *B* occasionally does not hold. In other words, if a statement can have exceptions, then that statement should not be asserted. Thus, propositional logic is intended to be applied only to statements that have no exceptions.

However, the rules in rule-based systems frequently have exceptions. Strictly speaking, one should not apply propositional logic to such rules.

But, might it not be permissible to ignore the occasional exception, to pretend that the rules contained in rule-based systems have no exceptions, and then to apply propositional logic?

This seems like it might be a reasonable procedure. Suppose that the rule *If A then B* has occasional exceptions but, nevertheless, we adopt the policy of applying this rule whenever we learn that *A* is true of an entity. Thus, whenever we learn that *A* is true, we will conclude that *B* is true. Most of the time when we conclude that *B* is true we will be correct. Occasionally, we will encounter an exception to the rule and our conclusion that *B* is true will be wrong. Shouldn't we be happy with this policy? Most of our conclusions are correct but an occasional conclusion is wrong. Isn't this better than a policy that never allowed us to reach any conclusion at all?

Putting the matter somewhat differently: If we employ rules that have exceptions, must we not be tolerant of rare errors in our conclusions? If we are not tolerant of rare errors, then we should not employ rules that have exceptions because they will inevitably cause us to occasionally reach incorrect conclusions.

Based upon the foregoing considerations, it might appear that, if we are tolerant of rare errors in our conclusions, then there are no further problems with applying propositional logic to rules with exceptions. Sadly, this is not the case.

2.2 A Worse Problem: Absurd Conclusions that Are Always Wrong

The policy of ignoring exceptions when applying propositional logic can result in conclusions that are more than just occasionally wrong. Certain of our conclusions will be absurd because they are always wrong. Here is an example.

We know that aircraft carriers are able to launch planes. This knowledge may be expressed in the form of a rule: *If aircraft carrier then able to launch planes*. This rule is denoted $A \Rightarrow L$, where A denotes the property of being an aircraft carrier, and L denotes the property of being able to launch planes. Since we are using propositional logic to reason with rules, the symbol " \Rightarrow " is interpreted as being the material-conditional connective.

But, of course, the rule *If aircraft carrier then able to launch planes* has exceptions, one being that aircraft carriers with broken catapults cannot launch planes. This exception may be written as a rule in its own right: $A \& B \Rightarrow \sim L$. Here B denotes the property of having a broken catapult.

Let D be the property of being a destroyer. Then, from $A \Rightarrow L$ and $A \& B \Rightarrow \sim L$, propositional logic derives the conclusion $A \& B \Rightarrow D$. In other

words, propositional logic concludes that aircraft carriers with broken catapults are destroyers.

An informal explanation of how propositional logic arrives at this conclusion is as follows. As mentioned, propositional logic is not designed to deal with statements that have exceptions. In effect, propositional logic takes the position that exceptions do not exist. Since carriers with broken catapults would be an exception to the rule that carriers are able to launch planes, propositional logic concludes that such carriers do not exist and, hence, that every one of them (there being none) is a destroyer.¹

It must be stressed that the above example does not show that there is something wrong with propositional logic. The problem is that we have used propositional logic to do something that it was not designed to do, namely, to reason with statements that can have exceptions. If we restrict our use of propositional logic to what it is designed to do, namely, to reason with statements that have no exceptions, then it works flawlessly.

In summary, when propositional logic is applied to rules that have exceptions, some of its conclusions are absurd (although they would be correct if the rules had no exceptions). Thus, propositional logic produces "too many" conclusions when applied to rules that have exceptions. We need a method of reasoning that produces fewer conclusions.

2.3 Censored Propositional Logic

Although propositional logic produces "too many" conclusions, it has continued to be employed in rule-based systems, but its conclusions have been "censored".

¹ There is nothing special about destroyers here. Using propositional logic, it may also be concluded that aircraft carriers with broken catapults are Marine battalions.

In diagnostic rule-based systems that reason by forward chaining [Winston, 1984, pp. 166–168; Parsaye and Chignell, 1988, pp. 271–273], this “censorship” is done as follows. At each stage of the system’s reasoning, the system possesses a collection of properties known (or, more accurately, believed) to hold true for the entity being diagnosed. Each of these properties was either told to the system or was deduced by the system. The system then checks to see which of its rules is applicable to the current collection of properties. (A rule *If A then B* is said to be *applicable* if *A* is contained in the current collection of properties.) Typically, at any stage of reasoning, multiple rules will be applicable. The system will select *one* of the applicable rules. The selected rule will then be applied. In other words, if the rule *If A then B* is applicable and has been selected, then the conclusion *B* will be deduced from the rule and then *B* will be added to the collection of properties.

The process of selecting one rule from multiple applicable rules is called *conflict resolution*. A variety of different strategies for conflict resolution have been used in rule-based systems [Winston, 1984, pp. 170–171; Parsaye and Chignell, 1988, pp. 273–275].

Conflict resolution is essentially a way of censoring the conclusions that could be drawn using propositional logic. Suppose that, at a given stage of reasoning, two rules are applicable: *If A then B* and *If C then D*. Since both rules are applicable, the current collection of properties contains both *A* and *C*. Hence, applying propositional logic, it would be legitimate to deduce both *B* and *D*. However, the process of conflict resolution will allow only one of the two rules to be applied. Hence, either *B* or *D* but not both will be added to the collection of properties. Thus, one conclusion deducible from propositional logic will be added to the collection of properties and the other will be censored. (However, the censorship may be only temporary. At a later stage of reasoning, the

previously censored rule may be selected and applied.)

In part, conflict resolution strategies are designed with the hope that they will avoid the deduction of irrelevant properties and, thus, will keep the inference process moving in a fruitful direction.

In addition, by censoring the conclusions that would otherwise be deduced, conflict resolution has the effect of reducing the likelihood that pairs of contradictory conclusions will be deduced. One strategy often used in conflict resolution is to apply the most specific rule. For example, suppose that the rules (i) *If aircraft carrier then able to launch planes* and (ii) *If aircraft carrier and has broken catapult then not able to launch planes* are both applicable at the same time. If both rules were applied, we would obtain a pair of contradictory conclusions: *able to launch planes* and *not able to launch planes*. However, the conflict resolution strategy based upon specificity will prevent both rules from being applied. Rule (ii) is considered to be more specific than Rule (i) because it is only applicable when Rule (i) is applicable but may be inapplicable when Rule (i) is applicable. Because Rule (ii) is more specific, it will be applied and Rule (i) will not be applied and, hence, a potential contradiction will be avoided.

As mentioned, a variety of conflict-resolution strategies have been used in rule-based systems. Moreover, the different strategies produce different results. So, which is the “correct” conflict-resolution strategy to use? Unfortunately, this question can’t be answered because there is no theory of conflict resolution. So, other than intuition or empirical results, we have no basis for preferring one conflict-resolution strategy over another.

It would be desirable to improve this situation. What we need in rule-based systems is a method

of reasoning that is based upon principles that are sound and well-understood.

3. Reinterpretation of Rules: Assertions of High Conditional Probability

We have seen that, because the rules used in rule-based systems typically have exceptions, those rules cannot be interpreted as statements in propositional logic. Specifically, when we interpret each rule *If A then B* (symbolized $A \Rightarrow B$) as being equivalent to the material conditional *Either not A or B* (symbolized $\sim A \vee B$), then unreasonable and/or contradictory conclusions may be deduced. This problem arises because we have pretended that the rules don't have exceptions when, in fact, they do.

This suggests that we should give up the pretense that the rules don't have exceptions. Instead we should interpret the rules in a way that frankly acknowledges the possibility of exceptions.

One way to acknowledge the possibility of exceptions would be to stop interpreting the symbol " \Rightarrow " as denoting the material conditional connective. Instead, we would give each rule $A \Rightarrow B$ an interpretation that cannot be expressed within propositional logic. Specifically, a rule $A \Rightarrow B$ would be interpreted as an *assertion of high conditional probability*, that is, as meaning that the conditional probability of B given A is close to one.² In other words, when A is true, B will be true *nearly* all the time, but perhaps not always; thus there may be rare exceptions where A is true but B is not.

² What is meant by "close to one"? This question cannot be given a precise answer; it depends upon our tolerance for error. As a rule of thumb: If exceptions to a rule $A \Rightarrow B$ are sufficiently rare that someone is willing to include that rule in a traditional rule-based system, then evidently the conditional probability of B given A must be "close to one" in their judgment.

Terminology. Let a rule $A \Rightarrow B$ be said to be *accurate* if the conditional probability of B given A is indeed close to one.

3.1 How to Do Diagnosis Under the New Interpretation of Rules

Suppose that we want to perform a diagnosis. Imagine that A_1, \dots, A_n summarize the known properties of the entity to be diagnosed. What we want to ascertain is whether B also holds true. So, the question we are faced with is: Given that A_1, \dots, A_n summarize the known properties of the entity, are we justified in concluding that that entity also possesses property B ?

Presumably, we are tolerant of rare errors in our conclusions because, if we were not tolerant of rare errors, then we should not employ rules that have exceptions. On the other hand, we are presumably not tolerant of frequent errors in our conclusions because, if we were tolerant of frequent errors, then we might as well answer questions by flipping a coin.

Given these tolerances for error, if A_1, \dots, A_n summarize the known properties of the entity to be diagnosed, then we should conclude that the entity has property B if and only if we believe that the conditional probability of B given A_1, \dots, A_n is close to one. This is equivalent to saying that we should infer the property B from the summary of known properties A_1, \dots, A_n if and only if we believe that $(A_1 \& \dots \& A_n) \Rightarrow B$ is an accurate rule.

When ought we to believe that $(A_1 \& \dots \& A_n) \Rightarrow B$ is an accurate rule? Presumably, if we are constructing a rule-based system, then we have in our possession a collection of all the rules that are empirically known to be accurate. So, obviously, if we find $(A_1 \& \dots \& A_n) \Rightarrow B$ in our collection of empirically known rules, then we should believe it to be accurate. More generally, if we can deduce $(A_1 \& \dots \& A_n) \Rightarrow B$ from the rules in our collection, then we should believe it to be accurate. However,

if we cannot deduce $(A_1 \& \dots \& A_n) \Rightarrow B$ from our collection of empirically known rules, then there is no reason why we should believe it to be accurate.

In summary, then, we should infer the property B from the summary of known properties A_1, \dots, A_n if and only if we can infer the rule $(A_1 \& \dots \& A_n) \Rightarrow B$ from the collection of empirically known rules.

3.2 Two Types of Reasoning Involved in Diagnosis

The above discussion shows that, to do diagnosis, we need to do two distinct types of reasoning: *property-rule-property reasoning* and *rule-rule reasoning*. Property-rule-property reasoning consists of inferring a property B from a summary of known properties A_1, \dots, A_n together with the rule $(A_1 \& \dots \& A_n) \Rightarrow B$. Rule-rule reasoning consists of inferring a rule from a collection of known rules.

3.2.1 Property-Rule-Property Reasoning

Because rules may have exceptions, the conclusions reached in property-rule-property reasoning are not guaranteed to be correct.

Property-rule-property reasoning may be said to be *nonmonotonic in the properties*. Loosely speaking, this means that learning new properties of an entity may cause us to withdraw previous conclusions about that entity. Here is an example. Suppose that the known rules are $A \Rightarrow L$ and $A \& B \Rightarrow \sim L$. If A is the only known property of an entity, then we may apply the rule $A \Rightarrow L$ to conclude that the entity has property L . On the other hand, if A and B are the known properties of the entity, then we may *not* apply the rule $A \Rightarrow L$ because A does not summarize all of the entity's known properties. Instead, we must apply the rule $A \& B \Rightarrow \sim L$ and conclude that the entity does *not* have the property L . Thus, if A is the only known property of the entity, we conclude that the entity has property L . (This conclusion is

tentative—rather than known for certain—because the rule $A \Rightarrow L$ may have exceptions.) If we subsequently learn that the entity also has the property B , then we withdraw our tentative conclusion that the entity has property L .

We shall shortly turn to a discussion of logics for rule-rule reasoning. But, before we do that, we will consider what rule-based systems for situation assessment should look like.

3.3 Rule-Based Systems for Situation Assessment

When rules are interpreted as being assertions of high conditional probability, then a rule-based system for situation assessment could be constructed as follows.

The system would have an entity-property base and a rule base. The entity-property base would contain information about each of the entities involved in the current situation; this information would come in the form of entity-property pairs. For example, in one pair, the property might be *aircraft*; this would indicate that the entity was an aircraft. The entity-property base would change continually as the situation changed. (Consideration of the mechanisms by which entity-property pairs are placed in the entity base and removed therefrom is beyond the scope of this paper.)

The system's rule base would contain all the known rules about entity properties. The contents of the rule base would be essentially constant, changing only with software upgrades.

The system would carry out deductions in the following manner. Suppose that it is desired to ascertain whether a particular entity, which will be termed the *entity of interest*, is a destroyer. (Let the property of being a destroyer be denoted D .) Then all the entity-property pairs involving the entity of interest would be retrieved from the

entity-property base. Suppose that is found that the properties possessed by the entity of interest are A_1, \dots, A_n . If the rule $(A_1 \& \dots \& A_n) \Rightarrow D$ is deduced from the rules in the rule base, then it is concluded that the entity of interest is a destroyer. On the other hand, if the rule $(A_1 \& \dots \& A_n) \Rightarrow \sim D$ is deduced from the rule base, then it is concluded that the entity of interest is *not* a destroyer. If neither rule can be deduced from the rule base, then no conclusion is reached concerning whether or not the entity of interest is a destroyer.

As is evident from the above discussion, a logic for deducing new rules from already-known rules is the key to rule-based situation assessment.

4. Deducing Rules from Rules

We need a logic for deducing rules from rules. Because we are interpreting rules as being assertions of high conditional probability, that means that we need a logic of high conditional-probability assertions.

4.1 Adams' Approach to Rule-Rule Reasoning

Ernest Adams [Adams, 1966, 1975] developed a logic of conditional statements. This logic may also be interpreted as a logic of high conditional-probability assertions [Adams, 1986, 1998]. So, because we are interpreting rules as being assertions of high conditional probability, Adams' logic may be used for rule-rule reasoning.

In Adams' logic, a rule $A \Rightarrow L$ is interpreted as meaning that the conditional probability of L given A is close to one. Thus, property A is almost always accompanied by the property L but exceptions are possible. In this logic, a *model* is defined to be a probability measure over the language used to express properties of entities.

Adams defined what it meant for a set of rules to be consistent. A loose intuitive explanation of Adams' formal definition may be given as follows.

A model (*i.e.*, probability measure) *satisfies* a rule $A \Rightarrow L$ if and only if the conditional probability of L given A is "close to one". A set of rules is *consistent* if, no matter how stringently we interpret "close to one", there exists a model that satisfies every rule in the set.³ After formulating his definition of consistency, Adams showed how to ascertain whether a set of rules is consistent.

Adams also defined what it meant for a rule to follow from (*i.e.*, be implied by) a set of rules. In formulating this definition, Adams used a criterion for reasoning that Bamber later named the *Criterion of Surety* [Bamber, to appear]. A loose intuitive explanation of this criterion is as follows. A conclusion rule *follows* from a set of premise rules if and only if *every* model that satisfies the premises also satisfies the conclusion. Equivalently, a conclusion rule $C \Rightarrow D$ follows from the premise rules $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n$ if and only if the conditional probabilities $\Pr(B_1|A_1), \dots, \Pr(B_n|A_n)$ all being close to one guarantees that the conditional probability $\Pr(D|C)$ will also be close to one.⁴ After formulating his definition of implication, Adams showed how to ascertain whether a conclusion rule is implied by a collection of premise rules.

4.1.1 Disturbing Results

Adams' results are disturbing in that they cast doubt on reasoning procedures used in rule-based systems. For example, suppose that it is desired to diagnose an entity having properties A and P . Suppose also that the rules $A \Rightarrow B$ and $P \Rightarrow Q$ are in the rule base. (Assume that A, B, P , and Q are distinct primitive properties.) In rule-based systems, it is allowable (unless prevented by the system's conflict resolution strategy) to apply the

³ To be precise: $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n$ are consistent if and only if, for every $\epsilon > 0$, there exists a probability measure \Pr such that $\Pr(B_1|A_1), \dots, \Pr(B_n|A_n)$ are all at least $1-\epsilon$.

⁴ To be precise: $C \Rightarrow D$ follows from $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n$ if and only if, for every $\epsilon > 0$, there exists a $\delta > 0$ such that, for all probability measures \Pr , if $\Pr(B_1|A_1), \dots, \Pr(B_n|A_n)$ are all at least $1-\delta$, then $\Pr(D|C)$ is at least $1-\epsilon$.

first rule to conclude that the entity has property B and later to apply the second rule to conclude that the entity has property Q . However, Adams' results show that, according to the Criterion of Surety, it is not valid to use the rules $A \Rightarrow B$ and $P \Rightarrow Q$ to conclude $A \& P \Rightarrow B \& Q$. (This is because it is possible for the conditional probability $\Pr(B \& Q | A \& P)$ to be far from one even though the conditional probabilities $\Pr(B | A)$ and $\Pr(Q | P)$ are both close to one.) In other words, the rule $A \& P \Rightarrow B \& Q$ can be inaccurate even though the rules $A \Rightarrow B$ and $P \Rightarrow Q$ are both accurate. But, if the rule $A \& P \Rightarrow B \& Q$ is inaccurate, then there is no justification for concluding that an entity having properties A and P should also have properties B and Q .

4.2 Bamber's Approach to Rule-Rule Reasoning

One possible reaction to the above example is to say that, although it may be theoretically possible for the rule $A \& P \Rightarrow B \& Q$ to be inaccurate when the rules $A \Rightarrow B$ and $P \Rightarrow Q$ are accurate, surely such situations are highly unusual. Moreover, if we are tolerant of rare errors in property-rule-property reasoning, then shouldn't we also be tolerant of rare errors in rule-rule reasoning? If so, then we ought to be willing to infer the rule $A \& P \Rightarrow B \& Q$ from the rules $A \Rightarrow B$ and $P \Rightarrow Q$ even though that inference may sometimes be in error.

This is not an unreasonable position. The difficulty with it, of course, is that the errors hypothesized to be rare (and, thus, tolerable) have not been demonstrated to actually be rare.

Bamber sought to solve this problem [Bamber, 1996; Bamber, to appear]. His work aimed at developing a logic for rule-rule reasoning in which rules derived as conclusions would be accurate nearly all the time but might be inaccurate on rare occasions. As previously, a *model* was defined to be a probability measure over the language used to express properties of entities. In addition, a second-order probability measure over models

was defined.⁵ In Bamber's approach, rather than deriving conclusions in accordance with the Criterion of Surety, conclusions were derived in accordance with the *Criterion of Near Surety*. Loosely put, this criterion states that a conclusion rule follows from a set of premise rules if and only if *nearly every* model that satisfies the premises also satisfies the conclusion.⁶ Equivalently, the conclusion rule $C \Rightarrow D$ follows from the premise rules $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n$ if and only if $\Pr(D | C)$ is close to one in *nearly every* model in which $\Pr(B_1 | A_1), \dots, \Pr(B_n | A_n)$ are close to one.⁷ Thus, the goal is not to infer conclusion rules that are *sure* to be accurate but rather to infer conclusion rules that are *nearly sure* to be accurate.

In summary, Bamber proposed a logic for rule-rule reasoning that was identical to Adams' logic except that the Criterion of Surety was replaced by the Criterion of Near Surety.

4.2.1 Comparison of the Two Criteria for Rule-Rule Reasoning

Recall the example that motivated consideration of the Criterion of Near Surety. It was found that, in Adams' logic (which employs the Criterion of Surety), the rule $A \& P \Rightarrow B \& Q$ does *not* follow from the rules $A \Rightarrow B$ and $P \Rightarrow Q$. However when

⁵ This was done as follows. Suppose that the language used to express properties contains k primitive properties. Hence, this language has 2^k atoms. Hence, any model may be represented by a 2^k -dimensional vector that specifies the probability of each atom. The second-order probability measure was defined to be the uniform distribution over all model vectors.

⁶ Bacchus *et al.* [Bacchus *et al.*, 1992, 1996; Grove *et al.*, 1994] investigated the Criterion of Near Surety using different types of models: (a) A *random-world model* was a first-order predicate calculus model. (b) A *random-structure model* was a collection of random-world models that were equivalent under permutation of domain elements. Bacchus *et al.* investigated the former approach much more deeply than the latter. The latter approach looks like it might yield results similar to those of Bamber, but this matter has not been definitively investigated.

⁷ To be precise: Let \Pr_{random} denote a random model (*i.e.*, a model selected at random from the second-order probability distribution). Then $C \Rightarrow D$ follows from $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n$ if and only if, for every $p > 0$ and every $e > 0$, there exists a $d > 0$ such that, $1-p$ is less than or equal to the conditional second-order probability that $\Pr_{\text{random}}(D | C)$ is at least $1-e$ given that $\Pr_{\text{random}}(B_1 | A_1), \dots, \Pr_{\text{random}}(B_n | A_n)$ are all at least $1-d$.

the Criterion of Near Surety is substituted for the Criterion of Surety, it is found that the rule $A \& P \Rightarrow B \& Q$ does indeed follow from the rules $A \Rightarrow B$ and $P \Rightarrow Q$.

Adams' logic for rule-rule reasoning may be said to be *monotonic in the rules*. This means that any conclusion that follows from a set of premise rules will still follow if the set of premise rules is enlarged. That is: If the conclusion $C \Rightarrow D$ follows from the premises $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n$, then $C \Rightarrow D$ will also follow from the enlarged set of premises $A_1 \Rightarrow B_1, \dots, A_n \Rightarrow B_n, A_{n+1} \Rightarrow B_{n+1}$.

On the other hand, if the Criterion of Near Surety is substituted for the Criterion of Surety, then the resulting logic is *not* monotonic in the rules. Here is a *trivial* example. Suppose that A , B , and C are primitive properties. Then the conclusion $A \& B \Rightarrow C$ follows from the premise $A \Rightarrow C$. Note however that it is possible for the conditional probability $\Pr(C|A)$ to be close to one and for the conditional probability $\Pr(C|A \& B)$ to be close to zero. Hence, the two rules $A \Rightarrow C$ and $A \& B \Rightarrow \sim C$ form a consistent set of premises. Obviously, from these premises, the conclusion $A \& B \Rightarrow C$ no longer follows although it did follow from the single premise $A \Rightarrow C$.

4.2.2 Choosing Between the Two Criteria for Rule-Rule Reasoning

Adams' and Bamber's approaches to rule-rule reasoning are essentially identical except that the former employs the Criterion of Surety whereas the latter employs the Criterion of Near Surety.

Which of these two criteria for rule-rule reasoning is more appropriate for use in rule-based systems? The Criterion of Near Surety yields all of the conclusions yielded by the Criterion of Surety plus some additional conclusions. However, we pay a price for these additional conclusions. Although the additional conclusions will nearly always be accurate, they will be inaccurate on rare occasions. So, in choosing

between the two reasoning criteria, we are choosing between fewer conclusions all of which are guaranteed to be accurate versus more conclusions some of which on rare occasions will be inaccurate. In short, we are faced with a trade-off.

Which side of this trade-off is preferred should depend upon our tolerance for rare errors. Evidently, in many rule-based systems, we do find it tolerable that the system makes rare errors. After all, if we weren't tolerant of rare errors, then we should never employ rules that have exceptions.

So, because we are tolerant of rule-based systems that produce rare errors, it seems reasonable that we should employ the Criterion of Near Surety in such systems.

Note that, in choosing between the two reasoning criteria, it is not a matter of deciding which criterion is "correct". Rather, it is a matter of deciding which criterion is more suitable for the task at hand.

4.3 Similarity of Bamber's Logic to Other Logics

Bamber showed [Bamber, to appear] that his logic is *formally identical*⁸ to the logic named Rational Closure [Lehmann and Magidor, 1992]. Bamber also showed that his logic has a close relationship to the logic named System-Z [Pearl, 1990]. The latter may be regarded as a bundling of a logic for property-rule-property reasoning together with a logic for rule-rule reasoning. Considered at a *formal* level, the rule-rule component of System-Z, is almost identical to

⁸ When two logics are said to be *formally identical*, it means that, as systems for reasoning with *uninterpreted* symbol strings, the two logics are identical. Thus, two logics can be formally identical even though the developers of the two logics interpret the meaning of the same symbol strings differently.

Bamber's logic.⁹ Efficient methods for carrying out deductions in these logics are known.

The formal identity/similarity between these three logics is interesting in that they were based upon different rationales. Thus, Bamber's logic was based upon the Criterion of Near Surety. On the other hand, System-Z may be interpreted as being based upon a principle of open-mindedness which may be loosely expressed: Don't doubt any possibility any more than the known rules compel one to do so [Bamber, to appear].

The fact that different rationales all result in the essentially the same logic is encouraging in that multiple rationales provide more convincing justification for using a logic than does a single rationale.

5. Summary

The following points were argued: Because the rules in diagnostic rule-based systems typically have exceptions, propositional logic is not an appropriate method for reasoning within diagnostic rule-based systems. Instead, rules should be interpreted probabilistically. That is, they should be regarded as assertions of high-conditional probability. Diagnostic reasoning involves (a) property-rule-property reasoning in which properties are inferred jointly from other properties and from rules and (b) rule-rule reasoning in which rules are inferred from rules. Rule-rule reasoning should be based upon the Criterion of Near Surety, which says that a conclusion rule follows from a set of premise rules if and only *nearly every* model of the premises is a model of the conclusion. In order to give a precise

⁹ The difference between the two logics has to do with a minor issue, namely, the consistency of rules involving impossible properties. For example, let U denote the property of being a unicorn and let W denote the property of being white. In System-Z, the rule $U \Rightarrow W \& \sim W$ is regarded as being self-inconsistent because nothing that exists can be both white and non-white. In contrast, in Bamber's logic, $U \Rightarrow W \& \sim W$ is regarded as being self-consistent and as expressing the concept that unicorns are non-existent.

formulation of this criterion, it is necessary to employ a second-order probability measure over models. The logic based upon the Criterion of Near Surety is formally similar or identical to logics based upon other rationales.

In the course of the paper, potential applications of this work were discussed. A scheme for constructing rule-based systems for situation assessment was expounded. Such situation-assessment systems would be useful either (a) as a decision aid for leaders of real-life military forces or (b) as a simulation of the situation assessment performed by computer generated forces.

References

[Adams, 1966] Ernest W. Adams. Probability and the logic of conditionals. In Jaakko Hintikka and Patrick Suppes (Editors), *Aspects of Inductive Logic*. Pp. 265–316. North Holland Publishing, Amsterdam, 1966.

[Adams, 1975] Ernest W. Adams. *The Logic of Conditionals: An Application of Probability to Deductive Logic*. D. Reidel Publishing, Dordrecht, Netherlands, 1975.

[Adams, 1986] Ernest W. Adams. On the logic of high probability. *Journal of Philosophical Logic*, 1986, 15, 255–279.

[Adams, 1998] Ernest W. Adams. *A Primer of Probability Logic*. Cambridge University Press, Cambridge, 1998.

[Bacchus *et al.*, 1992] Fahiem Bacchus, Adam J. Grove, Joseph Y. Halpern, and Daphne Koller. From statistics to belief. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*. 1992. Pp. 602–608.

[Bacchus *et al.*, 1996] Fahiem Bacchus, Adam J. Grove, Joseph Y. Halpern, and Daphne Koller.

From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 1996, **87**, 75–143.

[Bamber, 1996] Donald Bamber. Entailment in probability of thresholded generalizations. In Eric Horvitz and Finn Jensen (Editors), *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence (UAI-96)*. Pp. 57–64. Morgan Kaufmann Publishers, San Francisco, 1996.

[Bamber, to appear] Donald Bamber. Entailment with near surety of scaled assertions of high conditional probability. *Journal of Philosophical Logic*, to appear.

[Enderton, 1972] Herbert B. Enderton. *A Mathematical Introduction to Logic*. Academic Press, San Diego, 1972.

[Grove et al., 1994] Adam J. Grove, Joseph Y. Halpern, and Daphne Koller. Random worlds and maximum entropy. *Journal of Artificial Intelligence Research*, 1994, **2**, 33–88.

[Lehmann and Magidor, 1992] Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 1992, **55**, 1–60.

[Parsaye & Chignell, 1988] Kamran Parsaye and Mark Chignell. *Expert Systems for Experts*. Wiley, New York, 1988.

[Pearl, 1990] System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In Rohit Parikh (Editor), *Theoretical Aspects of Reasoning about Knowledge. Proceedings of the Third Conference (TARK 1990)*. Pp. 121–135. Morgan Kaufmann Publishers, San Mateo, CA.

[Winston, 1984] Patrick H. Winston. *Artificial Intelligence*, Second Edition. Addison-Wesley, Reading, MA, 1984.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1998		3. REPORT TYPE AND DATES COVERED Professional Paper	
4. TITLE AND SUBTITLE How Probability Theory Can Help Us Design Rule-Based Systems				5. FUNDING NUMBERS PR: CE31 PE: 0603707N WU: DN308298	
6. AUTHOR(S) D. Bamber				8. PERFORMING ORGANIZATION REPORT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Center San Diego, CA 92152-5001					
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Chief of Naval Operations Washington, DC 20350				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <p>This paper is concerned with finding improved methods of reasoning for use in rule-based systems that perform diagnosis in general and situation assessment is particular. It is argued that, because the rules used in rule-based systems typically have exceptions, the rules must be interpreted probabilistically. Thus, if a rule If A then B has exceptions, then what the rule really means is that the conditional probability of B given A is close to one. A rule-inference criterion that makes use of second-order probability concepts is advocated. Interestingly, this inference criterion is equivalent to some non-probabilistic inference criteria. The paper expounds a scheme for constructing situation-assessment systems that could be used as either decision aids or as reasoning components of computer generated forces.</p>					
Published in Proceedings of the 1998 Command and Control Research and Technology Symposium, pp. 441-451, June 1998.					
14. SUBJECT TERMS Mission Area: Command, Control, and Communications conflict resolution probability theory situation-assessment systems diagnostic rule-based system propositional logic logic rule-rule reasoning				15. NUMBER OF PAGES	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE		19. SECURITY CLASSIFICATION OF ABSTRACT	
				20. LIMITATION OF ABSTRACT Same as Report	

21a. NAME OF RESPONSIBLE INDIVIDUAL D. Bamber	21b. TELEPHONE <i>(include Area Code)</i> (619) 553-9219	21c. OFFICE SYMBOL D44215